



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|---|-------------|----------------------|-------------------------|------------------|
| 09/542,525 | 04/04/2000 | Gerald Dwayne Kuch | 3382-52327 | 5122 |
| 7590 | 03/25/2004 | | EXAMINER | |
| Klarquist Sparkman Campbell Leigh & Whinston LLP One World Trade Center Suite 1600 121 S W Salmon Street Portland, OR 97204 | | | NAHAR, QAMRUN | |
| | | | ART UNIT | PAPER NUMBER |
| | | | 2124 |) |
| | | | DATE MAILED: 03/25/2004 | |

Please find below and/or attached an Office communication concerning this application or proceeding.

| | | |
|------------------------------|-----------------|--------------|
| Office Action Summary | Application No. | Applicant(s) |
| | 09/542,525 | KUCH ET AL. |
| | Examiner | Art Unit |
| | Qamrun Nahar | 2124 |

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 04 April 2000.
 2a) This action is FINAL. 2b) This action is non-final.
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-53 is/are pending in the application.
 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
 5) Claim(s) _____ is/are allowed.
 6) Claim(s) 1-53 is/are rejected.
 7) Claim(s) _____ is/are objected to.
 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.
 10) The drawing(s) filed on 04 April 2000 is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
 Paper No(s)/Mail Date 4, 6, 8, 9, 10.

4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____.
 5) Notice of Informal Patent Application (PTO-152)
 6) Other: _____.

DETAILED ACTION

1. Claims 1-53 have been examined.

Information Disclosure Statement

2. The Information Disclosure Statement (IDS) filed on 3/20/01 is missing. Please submit another copy of the IDS filed on 3/20/01.

Applicant is advised that the re-submission of the IDS filed on 3/20/01 be accompanied by the acknowledgment postcard that indicates that the IDS was actually received by the Office.

Claim Objections

3. Claim 43 is objected to because of the following informalities: "The method of claim 42" on line 1 of the claim should be "The execution engine of claim 42". Appropriate correction is required.

Claim Rejections - 35 USC § 112

4. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

5. Claims 18-20, 27-30 and 47 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

6. Claim 18 recites the limitation "the null set" in line 16 of the claim. There is insufficient antecedent basis for this limitation in the claim. This limitation is interpreted as "a null set".

Claims 19-20 are rejected for dependency upon rejected base claim 18 above.

7. Claim 27 recites the limitation "the threshold" in line 14 of the claim. There is insufficient antecedent basis for this limitation in the claim. This limitation is interpreted as "the threshold framework".

Claims 28-30 are rejected for dependency upon rejected base claim 27 above.

8. Claim 47 recites the limitation "wherein **the** second entries" in line 3 of the claim. There is insufficient antecedent basis for this limitation in the claim. This limitation is interpreted as "wherein the entry for the second object class,".

Claim Rejections - 35 USC § 101

9. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

10. Claims 42-45 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

As per claim 42, merely claimed as a program representing a computer listing *per se* (execution engine), that is, descriptions or expressions of such a program and that is, descriptive material *per se*, non-functional descriptive material, and is not statutory because it is not a physical "thing" nor a statutory process, as there are not "acts" being performed. Such claimed programs do not define any structural and functional interrelationships between the program and

other claimed aspects of the invention which permit the program's functionality to be realized. Since a computer program is merely a set of instructions capable of being executed by a computer, the program itself is not a process, without the computer-readable medium needed to realize the program's functionality. In contrast, a claimed computer-readable medium encoded with a program defines structural and functional interrelationships between the program and the medium which permit the program's functionality to be realized, and is thus statutory.

Warmerdam, 33 F.3d at 1361, 31 USPQ2d at 1760. **In re Sarkar**, 588 F.2d 1330, 1333, 200 USPQ 132, 137 (CCPA 1978). See MPEP § 2106(IV)(B)(1)(a).

As per claims 43-45, these claims are rejected for dependency on the above rejected non-statutory claim 42.

Claim Rejections - 35 USC § 102

11. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

12. Claims 6-13, 15-21, 36, 39-51 and 53 are rejected under 35 U.S.C. 102(b) as being anticipated by Trishul M. Chilimbi and James R. Larus, “Using Generational Garbage Collection To Implement Cache-Conscious Data Placement”, 1998 (hereinafter “Chilimbi”).

Per Claim 6:

The Chilimbi publication discloses:

- in a computer, a method of arranging a plurality of data members of an object in a memory system (“This paper describes how a copying garbage collector can produce a cache-conscious object layout using real-time data profiling information. The copying phase of garbage collection offers a invaluable opportunity to reorganize a program’s data layout to improve cache performance. A cache-conscious data layout places objects with high temporal affinity near each other, so that they are likely to reside in the same cache block.” in pg. 37, col. 2, par. 5, lines 1-2 to pg. 38, col. 1, par. 1, lines 1-6)

- identifying more frequently referenced data members out of the plurality of data members of the object by consulting profiling data for the object; responsive to said identifying the more frequently referenced data members out of the plurality of data members of the object, grouping the more frequently referenced data members out of the plurality of data members of the object together into a group of more frequently referenced data members (“As described in Section 3, generational garbage collection copies live objects to *TO space*. Our goal is to use data profiling information to produce a cache-conscious layout of objects in *TO space* that places objects with high temporal affinity next to each other, so that they are likely to be in the same cache block. The data profiling information captures the temporal ordering of base object addresses, which our system uses to construct object affinity graphs. An object affinity graph is a weighted undirected graph in which nodes represent objects and edges encode temporal affinity between objects. ... The copying algorithm (Figure 9) described in this section uses the object affinity graph to produce a cache-conscious layout of

Art Unit: 2124

objects in *TO space*. ... STEP 1: Flip the roles of *FROM space* and *TO space*. Initialize the *unprocessed* and *free* pointers to the beginning of *TO space*. From the set of roots present in the affinity graph, pick the one with the highest affinity edge weight. Perform a greedy depth-first traversal of the entire object affinity graph starting from this node (i.e., visit the next unvisited node connected by the edge with greatest affinity weight). The stack depth for the depth-first traversal is limited to the number of nodes in the object affinity graph, and hence the object access buffer can be used as a scratch area for this purpose. In parallel with this greedy depth-first traversal, copy each object visited to *TO space* (increment the *free* pointer). Store this new object address as a forwarding address in the *FROM space* copy of the object.” in pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5)

- and responsive to said grouping the more frequently referenced data members into a group, assigning the data members from the group of more frequently referenced data members neighboring locations within the memory system (“After this step all affinity graph nodes will be laid out in *TO space* in a manner reflecting object affinities (Figure 9), but will still contain pointers to objects in *FROM space*.” in pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 7:

The Chilimbi publication discloses:

- wherein the memory system is a virtual memory system and the assigning assigns data members from the group of data members neighboring locations within the memory system by assigning data member from the group of data members a set of contiguous addresses in the virtual memory system (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 8:

The Chilimbi publication discloses:

- wherein said identifying more frequently referenced data members out of the plurality of data members of the object by consulting profiling data for the object comprises: ordering the data members according the number of memory references made during profiling of the object (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5).

Per Claim 9:

The Chilimbi publication discloses:

- wherein said identifying more frequently referenced data members out of the plurality of data members of the object by consulting profiling data for the object comprises: ordering the data members according to affinity between the data members (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5).

Per Claim 10:

The Chilimbi publication discloses:

- wherein said grouping the more frequently referenced data members out of the plurality of data members of the object together into a group of more frequently referenced data members comprises: computing an affinity matrix for the data members; transforming the affinity matrix into a cost matrix representing a graph of the data members; constructing a minimum-cost spanning tree for the graph; and traversing the minimum-cost spanning tree to determine members of the group (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 11:

The Chilimbi publication discloses:

- generating the profiling data by counting how many memory references were made to data members of the object during successive method calls (pg. 40, col. 2, par. 3, lines 1-5 to pg. 41, col. 1, par. 1, lines 1-7).

Per Claim 12:

The Chilimbi publication discloses:

- generating the profiling data by recording how many memory references were made to data members of the object during successive time periods (pg. 41, col. 2, par. 2, lines 1-11).

Per Claim 13:

The Chilimbi publication discloses:

- further comprising generating the profiling data in a database by a method comprising: maintaining counters indicative of how many memory references were made to data members of the object during successive observations; adding the counters to an overall counter; and writing a record to the database whenever the overall counter is likely to overflow based on an extrapolated forecast (pg. 40, col. 2, par. 3, lines 1-5 to pg. 41, col. 1, par. 1, lines 1-7 and par. 2, lines 1-9).

Per Claim 15:

The Chilimbi publication discloses:

- wherein the neighboring locations within the memory system are within a unit of the memory system separately loadable from other units of the memory system (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 16:

The Chilimbi publication discloses:

- wherein the unit of the memory system is a page separately loadable into primary memory (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 17:

The Chilimbi publication discloses:

- identifying less frequently referenced data members out of the plurality of data members of the object by consulting the profiling data for the object; responsive to said identifying the less frequently referenced data members out of the plurality of data members of the object, grouping the less frequently referenced data members out of the plurality of data members of the object together into a group of less frequently referenced data members (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5)

- and responsive to said grouping the less frequently referenced data members into a group, assigning the data members from the group of less frequently referenced data members at least one location separately-loadable from the locations of the more frequently referenced data members within the memory system (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 18:

The Chilimbi publication discloses:

- wherein the memory system is a virtual memory system having a primary memory and a secondary memory; the virtual memory system comprises memory units separately-loadable into the primary memory; said assigning the data members from the group of more frequently referenced data members neighboring locations within the memory system comprises assigning the more frequently referenced data members locations within a first set of one or more memory units in the virtual memory system; and said grouping the less frequently referenced data members out of the plurality of data members of the object together into a group of more frequently referenced data members comprises assigning the less frequently referenced data members locations within a second set of one or more memory units in the virtual memory system; and the intersection of the first set of memory units and the second set of memory units is a null set (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 19:

The Chilimbi publication discloses:

- wherein the memory units are memory pages (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 20:

The Chilimbi publication discloses:

- wherein the memory units are of variable size (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 21:

This is another version of the claimed method discussed above (claims 6 and 17), wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also anticipated by Chilimbi.

Per Claim 36:

This is another version of the claimed method discussed above, claim 21, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also anticipated by Chilimbi.

Per Claim 39:

The Chilimbi publication discloses:

- in a computer, a method of monitoring affinity among a plurality of data members of an object class in a memory system, the method comprising: constructing a table indicating how many times data members of the object were referenced during each of plurality of time periods during profiling (pg. 40, col. 2, par. 3, lines 1-5 to pg. 41, col. 1, par. 1, lines 1-7)

- computing an affinity matrix for the object class with a dot product operation, wherein the affinity matrix indicates an affinity between a first data member of the object and a

second data member of the object; and displaying the affinity matrix as a graph for consideration by a user (pg. 41, col. 1, par. 2, lines 1-9).

Per Claim 40:

The Chilimbi publication discloses:

- generating a database record by tracking how many time data members of the object are referenced during successive time periods by setting an expiring timer; and writing the database record to a database of profiling data upon expiration of the timer (pg. 41, col. 1, par. 2, lines 1-9).

Per Claim 41:

This is a system version of the claimed method discussed above, claim 21, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also anticipated by Chilimbi.

Per Claims 42 and 45:

These are execution engine versions of the claimed method discussed above, claim 21, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above, including “means for receiving a request from software to load the object of the object class” (pg. 40, col. 2, par. 3, lines 1-5 to pg. 41, col. 1, par. 1, lines 1-7). Thus, accordingly, these claims are also anticipated by Chilimbi.

Per Claim 43:

The Chilimbi publication discloses:

- wherein said means for arranging is a class loader operable to receive a layout directive specifying an explicit layout, a sequential layout, or an automatic layout (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 44:

The Chilimbi publication discloses:

- wherein the units of the virtual memory system are pages of a fixed size (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claims 46 and 49:

These are data structure versions of the claimed method discussed above, claim 21, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, these claims are also anticipated by Chilimbi.

Per Claim 47:

The Chilimbi publication discloses:

- in a set of data structures for indicating data layout, wherein the object class is a first object class and the data structure comprises an entry for a second object class, wherein the entry for the second object class, the layout directive field is defined to accommodate values for an automatic layout, an explicit layout, and a sequential layout (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 48:

The Chilimbi publication discloses:

- wherein the grouping field groups data members of the class according to their affinity as determined by calculating an affinity matrix for the data members and transforming the affinity matrix into a minimum-cost spanning tree (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 50:

This is a set of data structures version of the claimed data structure discussed above (claims 46 and 47), wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also anticipated by Chilimbi.

Per Claim 51:

The Chilimbi publication discloses:

- wherein the data structure for the respective object class further comprises: offset fields for the data members of the respective object class indicating at what position within a separately-loadable unit of the memory system the data members are to be placed (pg. 43, col. 2, par. 1, lines 2-8).

Per Claim 53:

This is another version of the claimed method discussed above, claim 21, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also anticipated by Chilimbi.

13. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

14. Claims 27, 29-30 and 34-35 are rejected under 35 U.S.C. 102(e) as being anticipated by Benitez (U.S. 6,189,141).

Per Claim 27:

The Benitez patent discloses:

- in a computer, a method of splitting data members of an object class into a set of groups to be placed at separately-loadable units of memory in a memory system by consulting profile data indicating how many times the data members were referenced in memory

(abstract and column 2, lines 35-52)

- assembling a list of the data members indicating how many times each data member in the list was referenced in memory; selecting a threshold value to be applied to the list of data members according to a threshold framework (column 2, lines 35-52)

- applying the threshold value to the list of data members to identify a candidate cold group; determining whether to split the data members based whether the candidate cold group is acceptable to a policy framework; and responsive to said determining, selecting between splitting the data members into a set of groups comprising a hot group and a cold group based on the threshold framework having affirmatively determined to split the data members and omitting to split the data members into a set of groups having determined not to split the data members (column 3, lines 5-17 and column 4, lines 20-56).

Per Claim 29:

The Benitez patent discloses:

- wherein the threshold framework supports specification of a default threshold value to be applied to a plurality of classes and specification of an override threshold value for a particular class (column 5, lines 10-18).

Per Claim 30:

The Benitez patent discloses:

- wherein the policy framework supports options based on a size of the candidate cold group and overhead is introduced into the cold group and a group referring to the cold group when the data members are split, the options comprising: determining whether to split the data members based on whether the size of the candidate cold group is greater than the overhead introduced into the referring group; determining whether to split the data members based on whether the size of the cold group is greater than a sum of the overhead introduced into the referring group and the overhead introduced into the cold group; determining whether to split the data members based on whether the size of the cold group is greater than a specified fraction of a sum of the overhead introduced into the referring group and the overhead introduced into the cold group; and determining whether to split the data members based on whether the size of the cold group is greater than a sum of the overhead introduced into the referring group and the overhead introduced into the cold group and whether the cold group is accessed more frequently than an average number of memory references (column 4, lines 20-56).

Per Claims 34 and 35:

These are another versions of the claimed method discussed above (claims 27 and 30), wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, these claims are also anticipated by Benitez.

Claim Rejections - 35 USC § 103

15. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

16. Claims 1-5, 14, 31-33, 37-38 and 52 are rejected under 35 U.S.C. 103(a) as being unpatentable over T. M. Chilimbi and James R. Larus, "Using Generational Garbage Collection To Implement Cache-Conscious Data Placement", 1998 (hereinafter "Chilimbi") in view of Li (U.S. 6,631,496).

Per Claim 1:

Chilimbi teaches in a computer, a method of arranging, a plurality of data members of an object class in a virtual memory system having units separately loadable from secondary memory into primary memory ("This paper describes how a copying garbage collector can produce a cache-conscious object layout using real-time data profiling information. The copying phase of garbage collection offers a invaluable opportunity to reorganize a program's data layout to improve cache performance. A cache-conscious data layout places objects with high temporal

affinity near each other, so that they are likely to reside in the same cache block.” in pg. 37, col. 2, par. 5, lines 1-2 to pg. 38, col. 1, par. 1, lines 1-6), the method comprising: consulting profiling data indicating how the data members of the object class are to be grouped into a plurality of separate groups comprising a first group and a second group (“As described in Section 3, generational garbage collection copies live objects to *TO space*. Our goal is to use data profiling information to produce a cache-conscious layout of objects in *TO space* that places objects with high temporal affinity next to each other, so that they are likely to be in the same cache block. The data profiling information captures the temporal ordering of base object addresses, which our system uses to construct object affinity graphs. An object affinity graph is a weighted undirected graph in which nodes represent objects and edges encode temporal affinity between objects. … The copying algorithm (Figure 9) described in this section uses the object affinity graph to produce a cache-conscious layout of objects in *TO space*. … STEP 1: Flip the roles of *FROM space* and *TO space*. Initialize the *unprocessed* and *free* pointers to the beginning of *TO space*. From the set of roots present in the affinity graph, pick the one with the highest affinity edge weight. Perform a greedy depth-first traversal of the entire object affinity graph starting from this node (i.e., visit the next unvisited node connected by the edge with greatest affinity weight). The stack depth for the depth-first traversal is limited to the number of nodes in the object affinity graph, and hence the object access buffer can be used as a scratch area for this purpose. In parallel with this greedy depth-first traversal, copy each object visited to *TO space* (increment the *free* pointer). Store this new object address as a forwarding address in the *FROM space* copy of the object.” in pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5); responsive to said

Art Unit: 2124

consulting profiling data indicating how the data members of the object class are to be grouped into a plurality of separate groups comprising a first group and a second group, assigning memory locations for data members of the first group within a first unit of memory in the virtual memory system; and responsive to said consulting profiling data indicating how the data members of the object are to be grouped into a plurality of separate groups comprising a first group and a second group, assigning memory locations for data members of the second group within a second unit of memory in the virtual memory system separately loadable into primary memory from the first unit (“After this step all affinity graph nodes will be laid out in *TO space* in a manner reflecting object affinities (Figure 9), but will still contain pointers to objects in *FROM space*.” in pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9). Chilimbi does not explicitly teach storing profiling data as metadata. Li teaches storing profiling data as metadata (column 6, lines 25-28).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by Chilimbi to include storing profiling data as metadata using the teaching of Li. The modification would be obvious because one of ordinary skill in the art would be motivated to ease the retrieval of profile data by storing it as metadata.

Per Claim 2:

The rejection of claim 1 is incorporated, and Chilimbi further teaches wherein the plurality of separate groups comprises more than two groups, the method further comprising: responsive to said consulting profiling data indicating how the data members of the object class

Art Unit: 2124

are to be grouped into a plurality of separate groups comprising more than two groups, assigning memory locations for data members of groups other than the first group and the second group within respective units of memory in the virtual memory system separately loadable from the first unit of memory and the second unit of memory (pg. 43, col. 2, par. 1, lines 6-8 and Fig. 9).

Per Claim 3:

The rejection of claim 1 is incorporated, and Chilimbi further teaches based on profile data of the object, generating an affinity tree to indicate more frequently referenced data members of the object are to be grouped in the first group and less frequently referenced data members are to be grouped in the second group (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5). Chilimbi does not explicitly teach generating metadata based on profile data. Li teaches generating metadata based on profile data (column 6, lines 25-28).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by Chilimbi to include generating metadata based on profile data using the teaching of Li. The modification would be obvious because one of ordinary skill in the art would be motivated to ease the retrieval of profile data by storing it as metadata.

Per Claim 4:

The rejection of claim 1 is incorporated, and Chilimbi further teaches based on profile data of the object, generating an affinity tree to indicate data members of the object with greater

Art Unit: 2124

affinity are to be grouped in the first group and data members with lesser affinity are to be grouped in the second group (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5). Chilimbi does not explicitly teach generating metadata based on profile data. Li teaches generating metadata based on profile data (column 6, lines 25-28).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by Chilimbi to include generating metadata based on profile data using the teaching of Li. The modification would be obvious because one of ordinary skill in the art would be motivated to ease the retrieval of profile data by storing it as metadata.

Per Claim 5:

The rejection of claim 4 is incorporated, and Chilimbi further teaches generating the profile data by recording how many memory references were made to data members of the object during successive substantially-equal time periods; and calculating an affinity matrix from the profiling data (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-11; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 14:

This is another version of the claimed method discussed above, claim 3, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also obvious.

Per Claims 31 and 33:

These are another versions of the claimed method discussed above (claims 1 and 3), wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, these claims are also obvious.

Per Claim 32:

The rejection of claim 31 is incorporated, and Chilimbi further teaches declaratively specifying in the software that a specified object class definition is not to be optimized; and during said grouping, inhibiting grouping of data members of the specified object class definition (pg. 40, col. 1, par. 2, lines 1-8 to col. 2, par. 2, lines 1-10).

Per Claim 37:

This is another version of the claimed method discussed above, claim 31, wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also obvious.

Per Claim 38:

The rejection of claim 37 is incorporated, and Chilimbi further teaches wherein one or more data members appear in the minimum-cost spanning tree in a graph disconnected from the connected graph, the method further comprising: responsive to determining the data members appear in a graph disconnected from the connected graph, placing data members in the

disconnected graph in a group separate from the hot group (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5).

Per Claim 52:

This is another version of the claimed method discussed above (claims 1 and 3), wherein all claim limitations also have been addressed and/or covered in cited areas as set forth above. Thus, accordingly, this claim is also obvious.

17. Claims 22 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over T. M. Chilimbi and James R. Larus, "Using Generational Garbage Collection To Implement Cache-Conscious Data Placement", 1998 (hereinafter "Chilimbi") in view of Benitez (U.S. 6,189,141).

Per Claim 22:

The rejection of claim 21 is incorporated, and Chilimbi further teaches wherein said grouping the more frequently referenced data members out of the plurality of data members of the object together into first group of data members and the less frequently referenced data members out of the plurality of data members of the object together into a second group of data members comprises the following: ordering the data members by how many times referenced in memory during profiling (pg. 41, col. 1, par. 3, lines 1-6 to col. 2, par. 1, lines 1-5; pg. 43, col. 1, par. 3, lines 5-8; pg. 43, col. 1, par. 5, lines 1-9 to col. 2, par. 1, lines 1-5). Chilimbi does not explicitly teach dividing the data members into the first group and the second group based on a

threshold. Benitez teaches dividing the data members into the first group and the second group based on a threshold (abstract and column 10, lines 43-53).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by Chilimbi to include dividing the data members into the first group and the second group based on a threshold using the teaching of Benitez. The modification would be obvious because one of ordinary skill in the art would be motivated to dynamically optimize a code block during each execution/trace pass.

Per Claim 25:

The rejection of claim 22 is incorporated, Benitez further teaches receiving a specified number of memory references indicative of a desired threshold number of memory references observed during profiling; wherein the threshold is the desired threshold number of memory references (column 2, lines 45-52).

18. Claim 28 is rejected under 35 U.S.C. 103(a) as being unpatentable over Benitez (U.S. 6,189,141) in view of Christenson (U.S. 6,324,620).

Per Claim 28:

The rejection of claim 27 is incorporated, and further, Benitez does not explicitly teach wherein the threshold framework supports options comprising: specifying a threshold value in terms of absolute number of memory references to data members of the object class; specifying a threshold value in terms of multiples of an average number of memory references to data

members of the object class; and specifying a threshold value in terms of standard deviations from an average number memory referenced to data members of the object class. Christenson teaches wherein the threshold framework supports options comprising: specifying a threshold value in terms of absolute number of memory references to data members of the object class; specifying a threshold value in terms of multiples of an average number of memory references to data members of the object class; and specifying a threshold value in terms of standard deviations from an average number memory referenced to data members of the object class (column 3, lines 20-45).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by Benitez to include that the threshold framework supports options comprising: specifying a threshold value in terms of absolute number of memory references to data members of the object class; specifying a threshold value in terms of multiples of an average number of memory references to data members of the object class; and specifying a threshold value in terms of standard deviations from an average number memory referenced to data members of the object class using the teaching of Christenson. The modification would be obvious because one of ordinary skill in the art would be motivated to reduce the constant movement of data in the memory of the computer system.

19. Claims 23, 24 and 26 are rejected under 35 U.S.C. 103(a) as being unpatentable over T. M. Chilimbi and James R. Larus, "Using Generational Garbage Collection To Implement Cache-Conscious Data Placement", 1998 (hereinafter "Chilimbi") in view of Benitez (U.S. 6,189,141), and further in view of Christenson (U.S. 6,324,620).

Per Claims 23, 24 & 26:

The rejection of claim 22 is incorporated, and further, the combination of Chilimbi and Benitez does not explicitly teach calculating an average number of references observed during profiling; wherein the threshold is the average number of references or calculating an average number of references observed during profiling; receiving a specified number representing a desired multiple of an average; wherein the threshold is the desired multiple times the average number of references or calculating a standard deviation for a number of references observed during profiling; receiving a specified number representing a desired multiple of the standard deviation; wherein the threshold is the desired multiple times the standard deviation.

Christenson teaches calculating an average number of references observed during profiling; wherein the threshold is the average number of references; calculating an average number of references observed during profiling; receiving a specified number representing a desired multiple of an average; wherein the threshold is the desired multiple times the average number of references; and calculating a standard deviation for a number of references observed during profiling; receiving a specified number representing a desired multiple of the standard deviation; wherein the threshold is the desired multiple times the standard deviation (column 3, lines 20-45).

It would have been obvious to one having ordinary skill in the computer art at the time of the invention was made to modify the method disclosed by the combination of Chilimbi and Benitez to include calculating an average number of references observed during profiling; wherein the threshold is the average number of references; calculating an average number of

references observed during profiling; receiving a specified number representing a desired multiple of an average; wherein the threshold is the desired multiple times the average number of references; and calculating a standard deviation for a number of references observed during profiling; receiving a specified number representing a desired multiple of the standard deviation; wherein the threshold is the desired multiple times the standard deviation using the teaching of Christenson. The modification would be obvious because one of ordinary skill in the art would be motivated to reduce the constant movement of data in the memory of the computer system.

Conclusion

20. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

21. Any inquiry concerning this communication from the examiner should be directed to Qamrun Nahar whose telephone number is (703) 305-7699. The examiner can normally be reached on Mondays through Thursdays from 9:00 AM to 6:30 PM. The examiner can also be reached on alternate Fridays.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kakali Chaki, can be reached on (703) 305-9662. The fax phone number for the organization where this application or processing is assigned is (703) 872-9306.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

Application/Control Number: 09/542,525
Art Unit: 2124

Page 30

QN
March 16, 2004



TODD INGERBERG
PRIMARY EXAMINER